


# Hands on with Service Mesh

NYRHUG December 2019

---

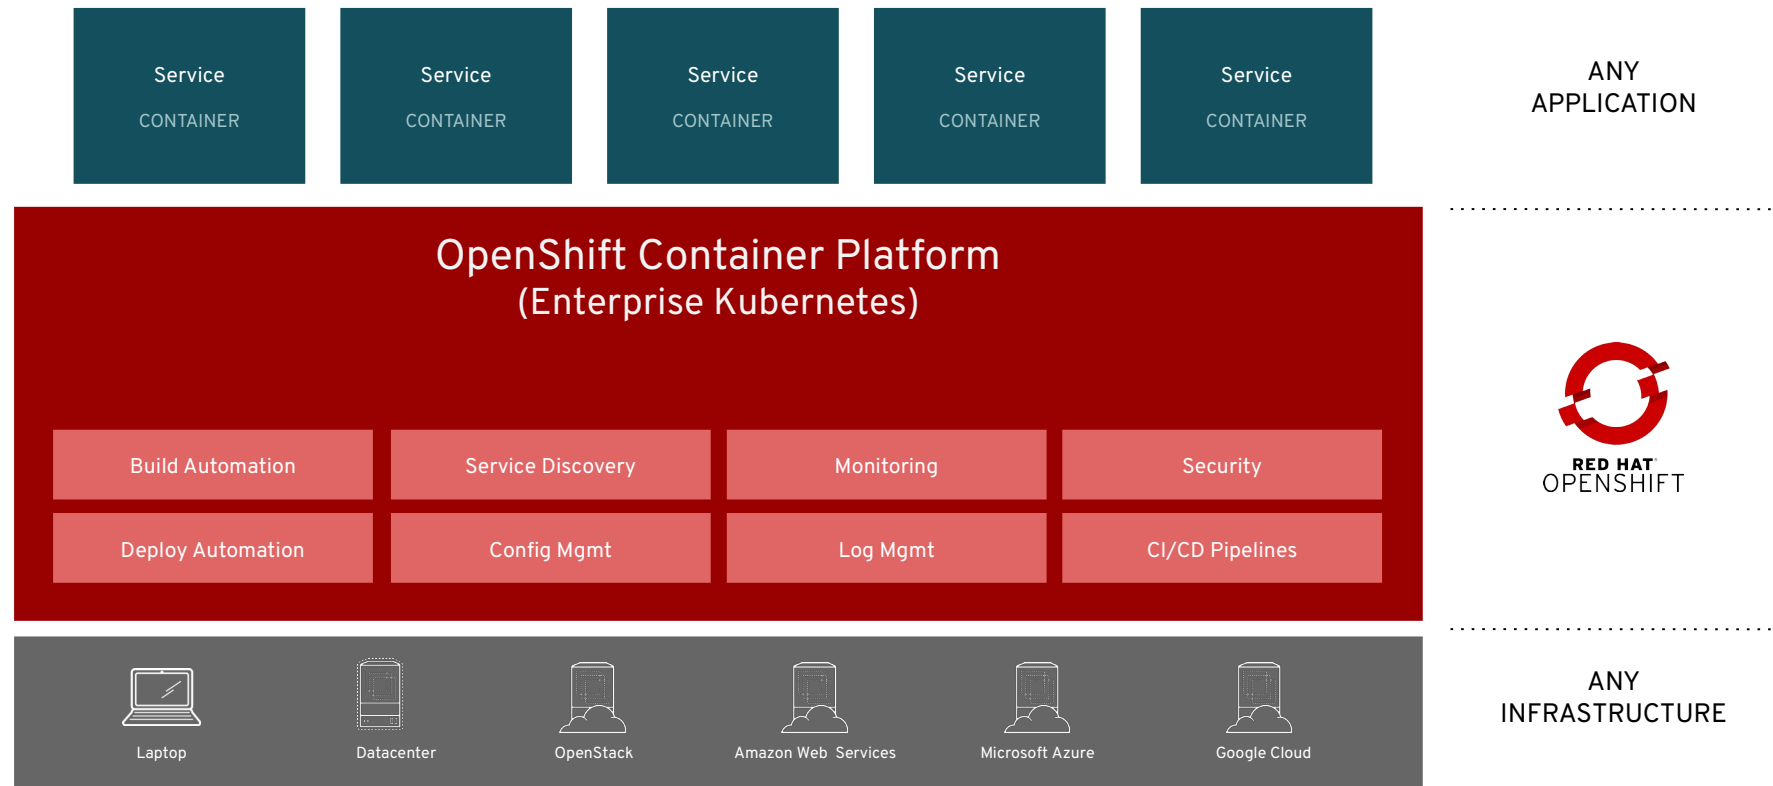
Patrick Ladd  
Technical Account Manager



A mash-up of several better-known technologies: “A service mesh is a set of software components which act as the “glue” for a set of independent applications. The goal of the mesh is to guarantee secure communications between each application and be able to redirect traffic in the event of failures. Often the features of a service mesh look like a mash-up between a load balancer, a web application firewall, and an API gateway.”

- Brian “Redbeard” Harrington, Product Manager at Red Hat

# BUILD AND DEPLOY CLOUD-NATIVE APPS WITH RED HAT OPENSSHIFT




# DEVELOPMENT PAIN POINTS

Click to add subtitle




## OUR SECURITY TEAM SAYS WE HAVE TO INTRODUCE \_\_\_\_\_

Security needs can change quickly. Many times this can require near-constant research to stay on top of the latest trends.



## I'VE NEVER RUN A PRODUCTION FACING SERVICE DIRECTLY

If one has never “carried the pager” for a production service, it can be challenging to foresee the needs to of those who do.




## DID THIS CODE CHANGE FIX MY PERFORMANCE ISSUE?

It can be challenging to ensure that bugs are actually squashed. We need better ways of measuring changes to applications, ideally in a deterministic manner.

# MANAGEMENT PAIN POINTS

Click to add subtitle




## IF I CAN'T SSH INTO THE HOST, I CAN'T DEBUG IT

Users have conflated *how* they achieve certain goals with the goal itself. There are bad behaviors which we must solve in other ways.



## AUDITING OF EXISTING RESOURCES CAN BE CHALLENGING

Do you know which services are safe to shut off, or how they are connected?




## WHERE ARE THE LOGS FOR THE APPLICATION?

As monoliths are broken into microservices there are many more places to search for logs. As these are scheduled across hosts as containers, this challenge grows exponentially.


# MANAGEMENT PAIN POINTS

Click to add subtitle




## I DON'T UNDERSTAND ALL OF THESE APPLICATION COMPONENTS

Gabriel Monroy (Founder of Deis, former CTO of Engine Yard) summarized it well “DevOps means developers writing Puppet manifests and SysAdmins sitting through architecture meetings for software they will never understand.”



## “WHY DON'T YOU JUST \_\_\_\_\_?”

Often, the needs of each side are hard to understand. Developers may not see the importance of on-the-fly reconfiguration while SREs may not fully appreciate the need for application level tracing in a constellation of microservices.




## I NEED TO CHANGE HOW TRAFFIC IS FLOWING

There are times when the easiest solution to a temporary problem is to redirect users and their traffic. If these features have not been built into an application it can be challenging (if not impossible) for operations to achieve this.


# DEPLOYMENT PAIN POINTS

Click to add subtitle




## THERE IS NO REASON XXX SHOULD BE ABLE TO HIT YYY

Without the sufficient runbooks, understanding the flow of an application can be obtuse. Without understanding the flow, mitigating emerging security threats can be impossible.



## IF I ROTATE THE CERTIFICATES\* WILL THE APPLICATION BREAK?

Fear often drives the decisions of SysAdmins. Without the ability to test the outcome of changes, they will default to the most safe process.



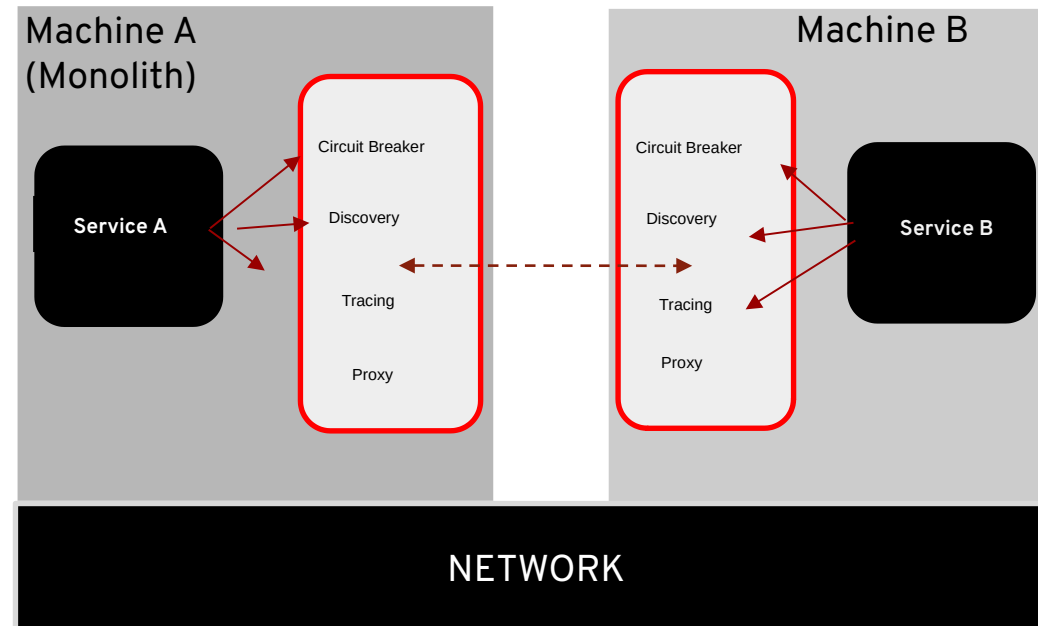
## I NEED TO TEST THIS NEW VERSION BEFORE DEPLOYING IT

A lack of real-time testing of software will lead to limited windows during which software can be deployed. Failures in this process lead to more draconian measures like [ITIL](#).

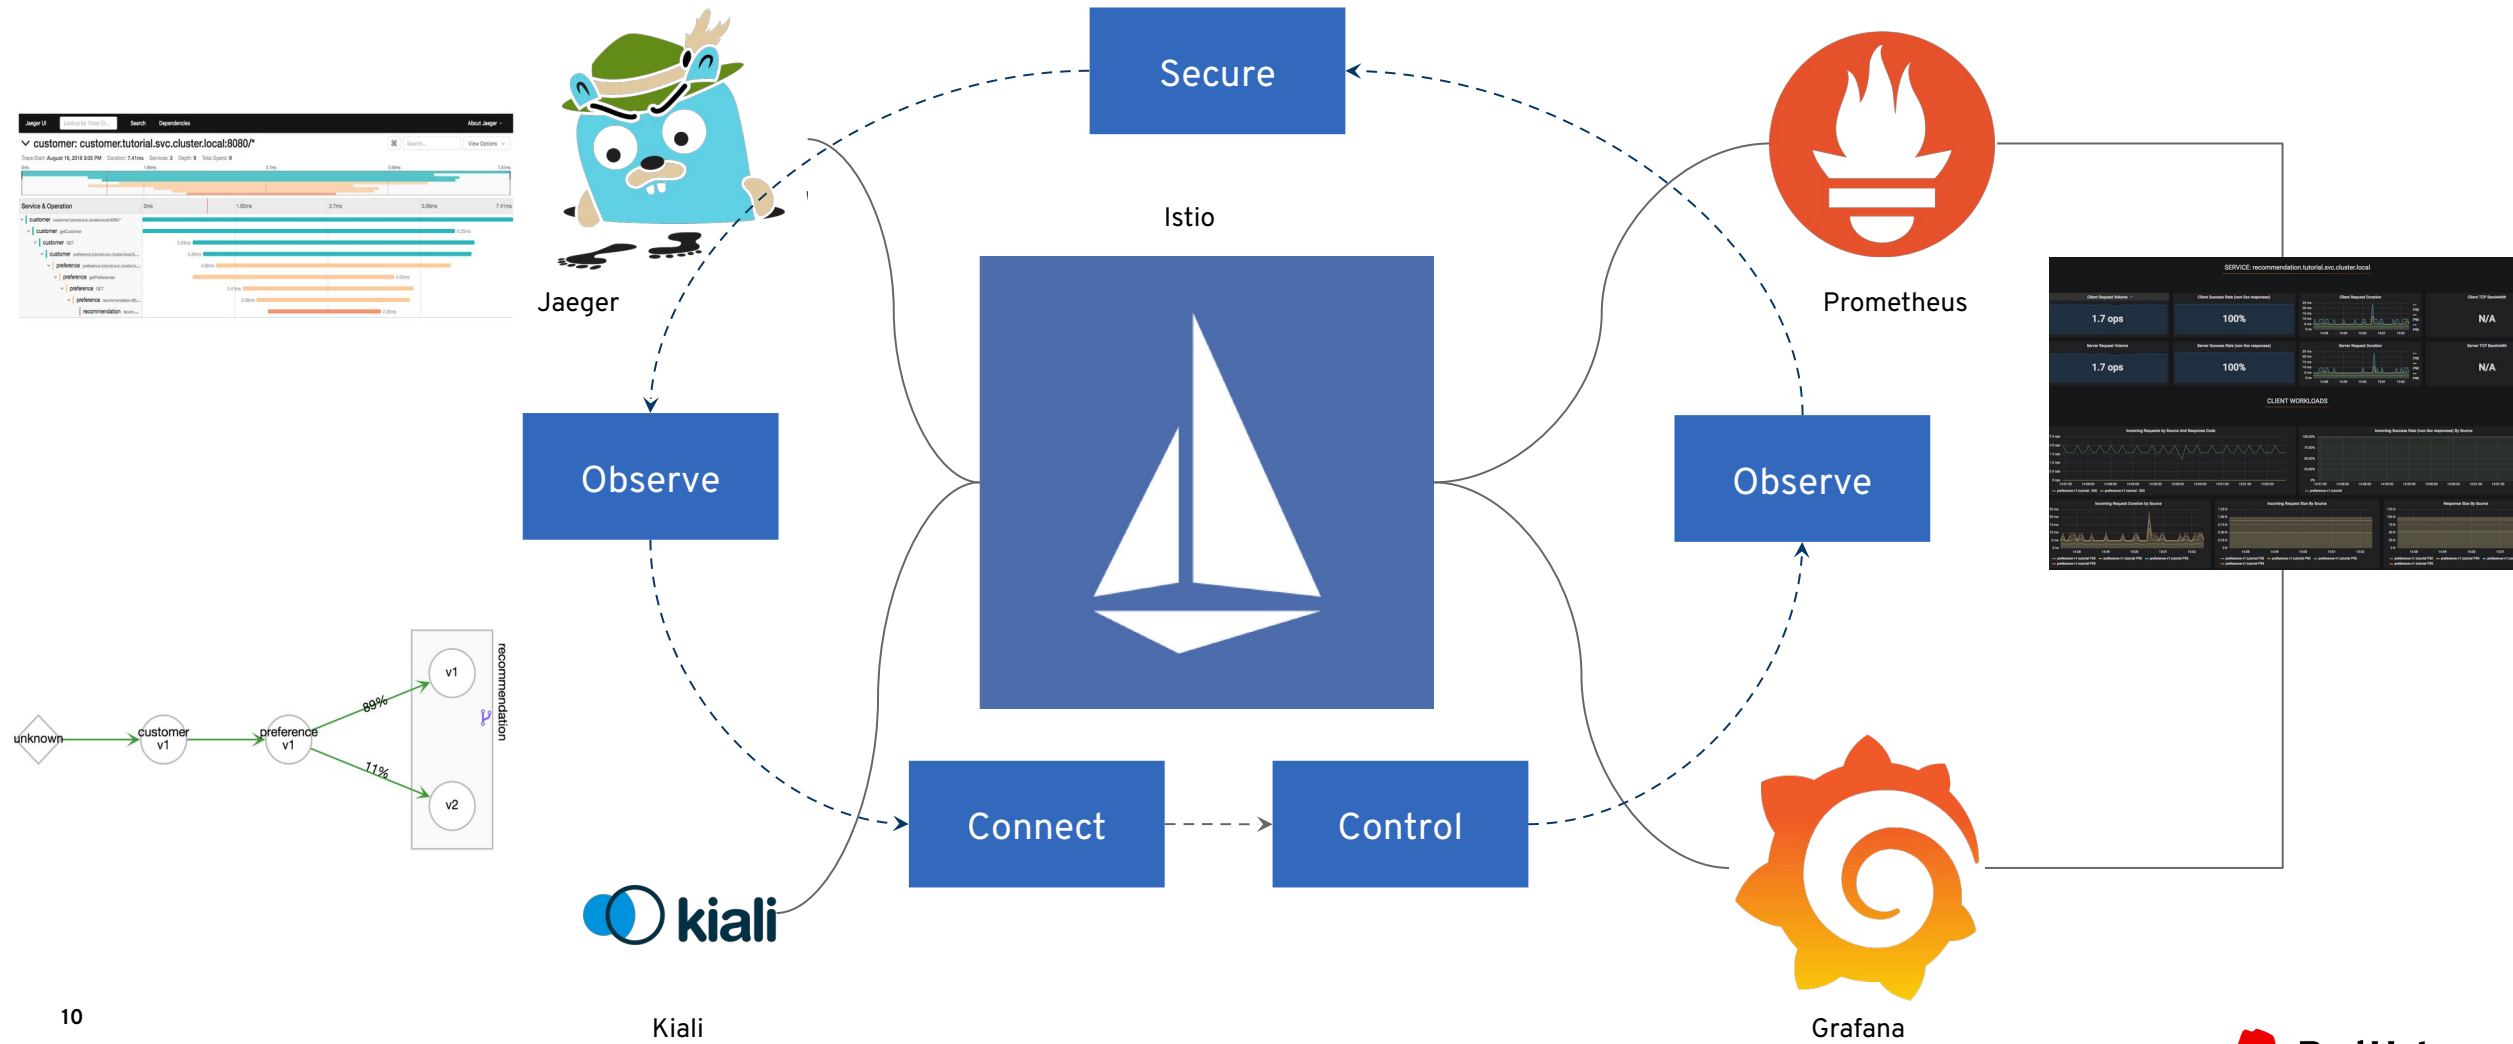
# OVERVIEW



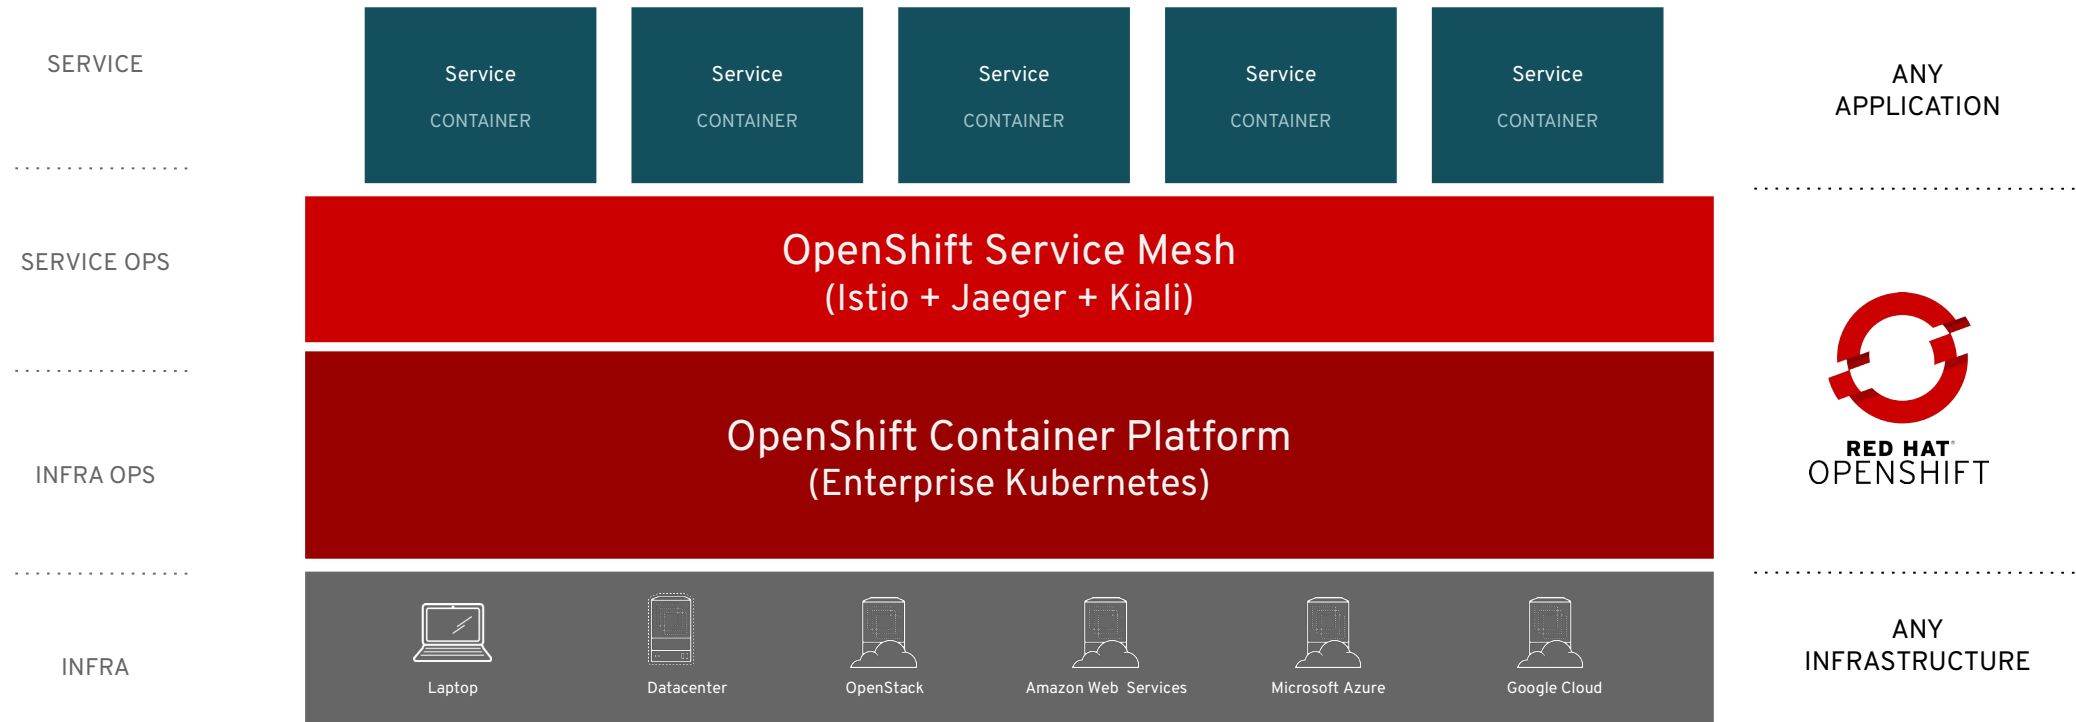
# WHAT IS A SERVICE MESH ?



# SERVICE MESH ECOSYSTEM

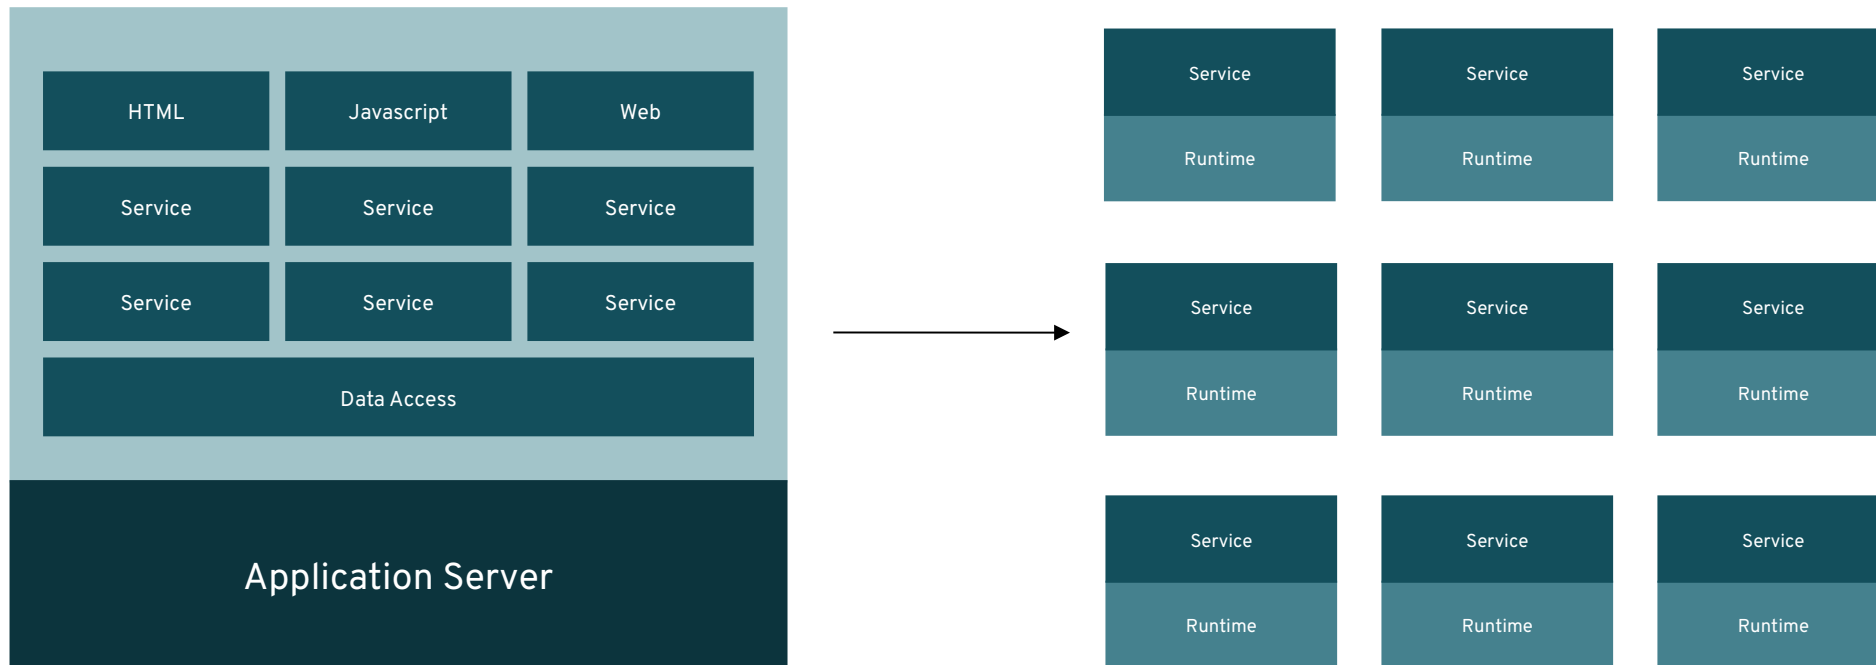


# DISTRIBUTED SERVICES WITH RED HAT OPENSIFT SERVICE MESH



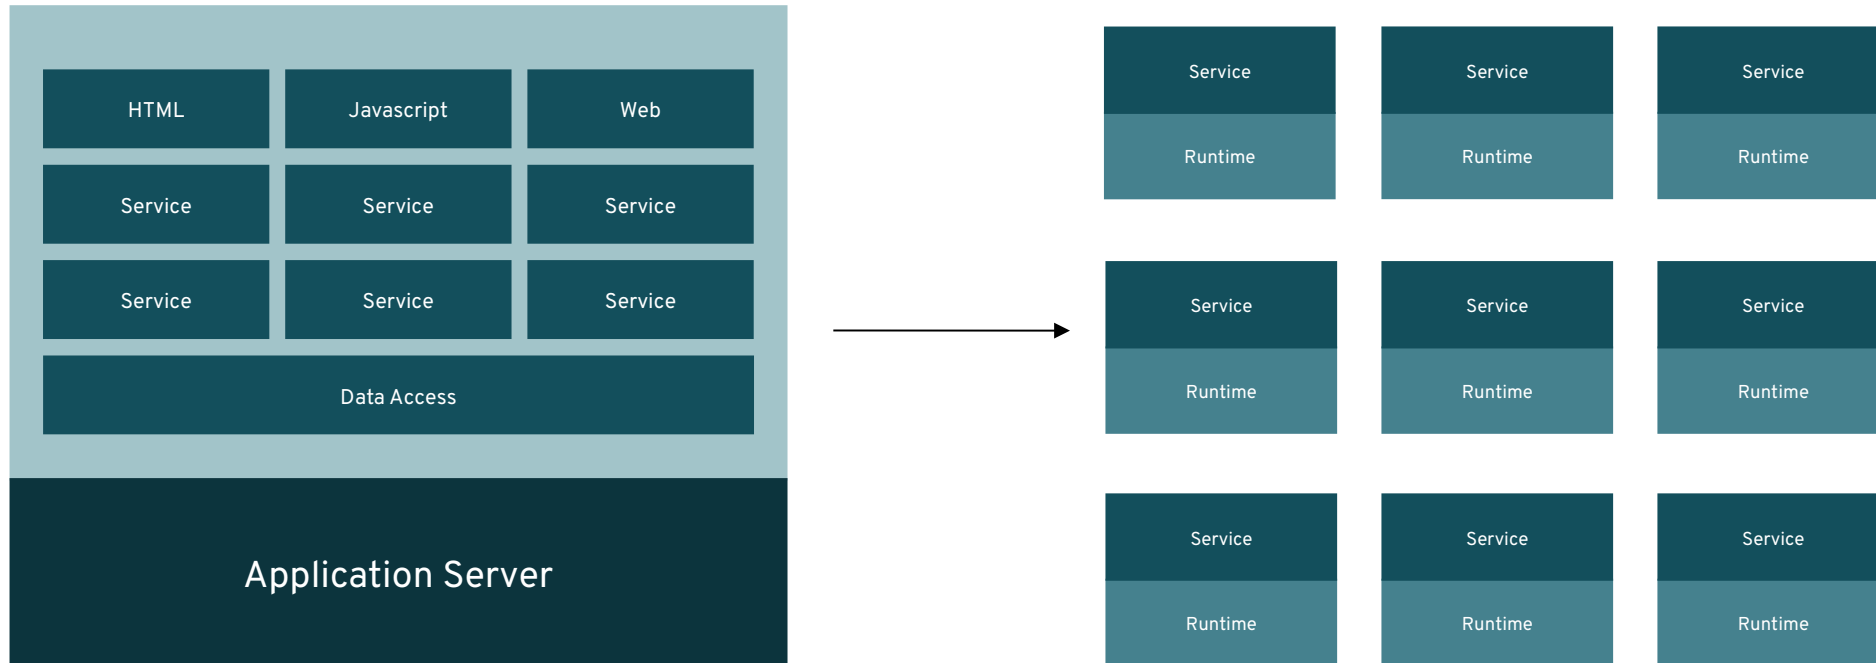
# UNDER THE HOOD

# MICROSERVICES ARCHITECTURE




# MICROSERVICES ARCHITECTURE

## DISTRIBUTED



# DISTRIBUTED ARCHITECTURE

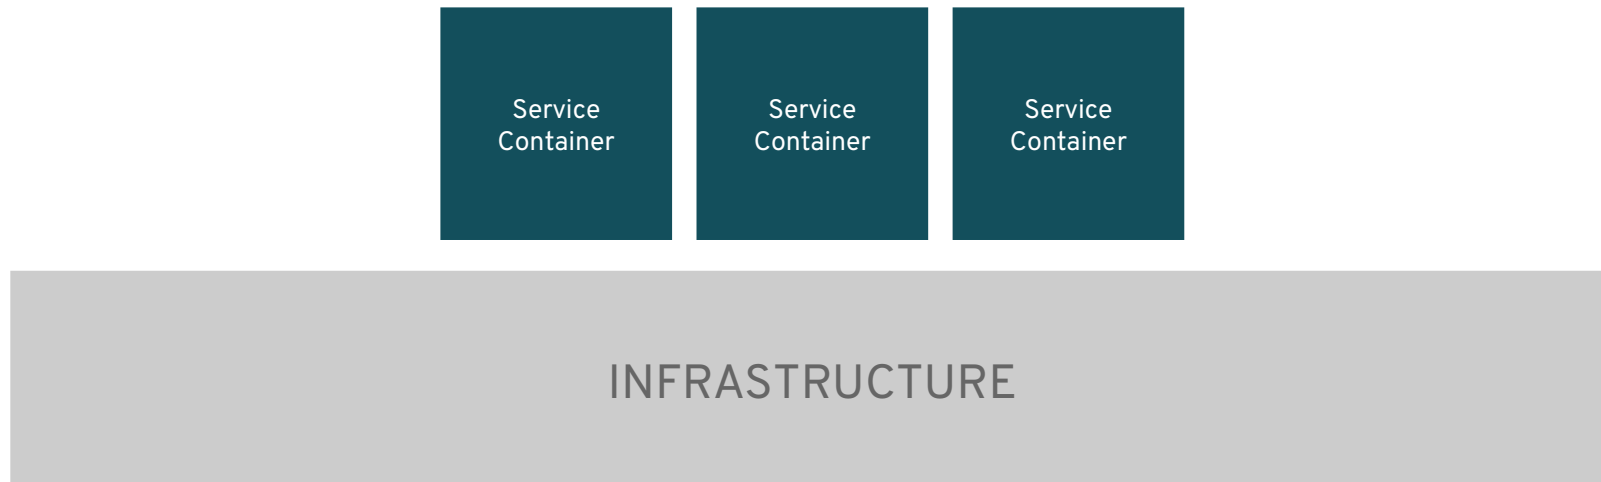








# HOW TO DEAL WITH THE COMPLEXITY?



# DEPLOYMENT

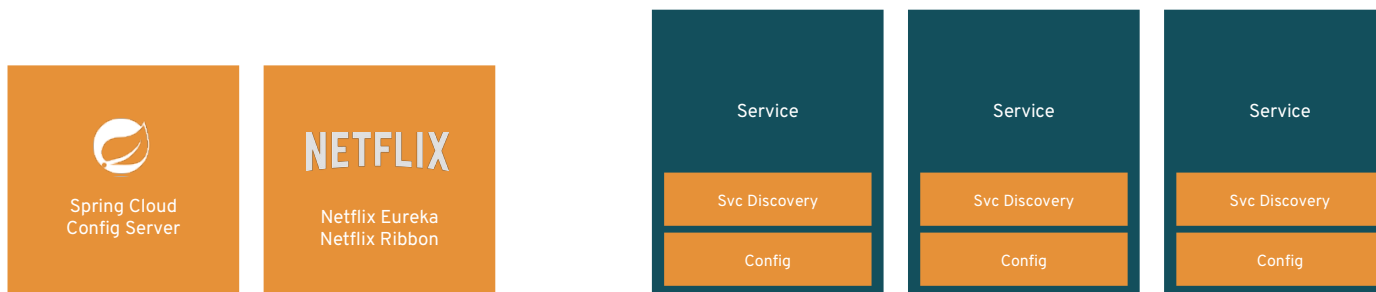


# CONFIGURATION



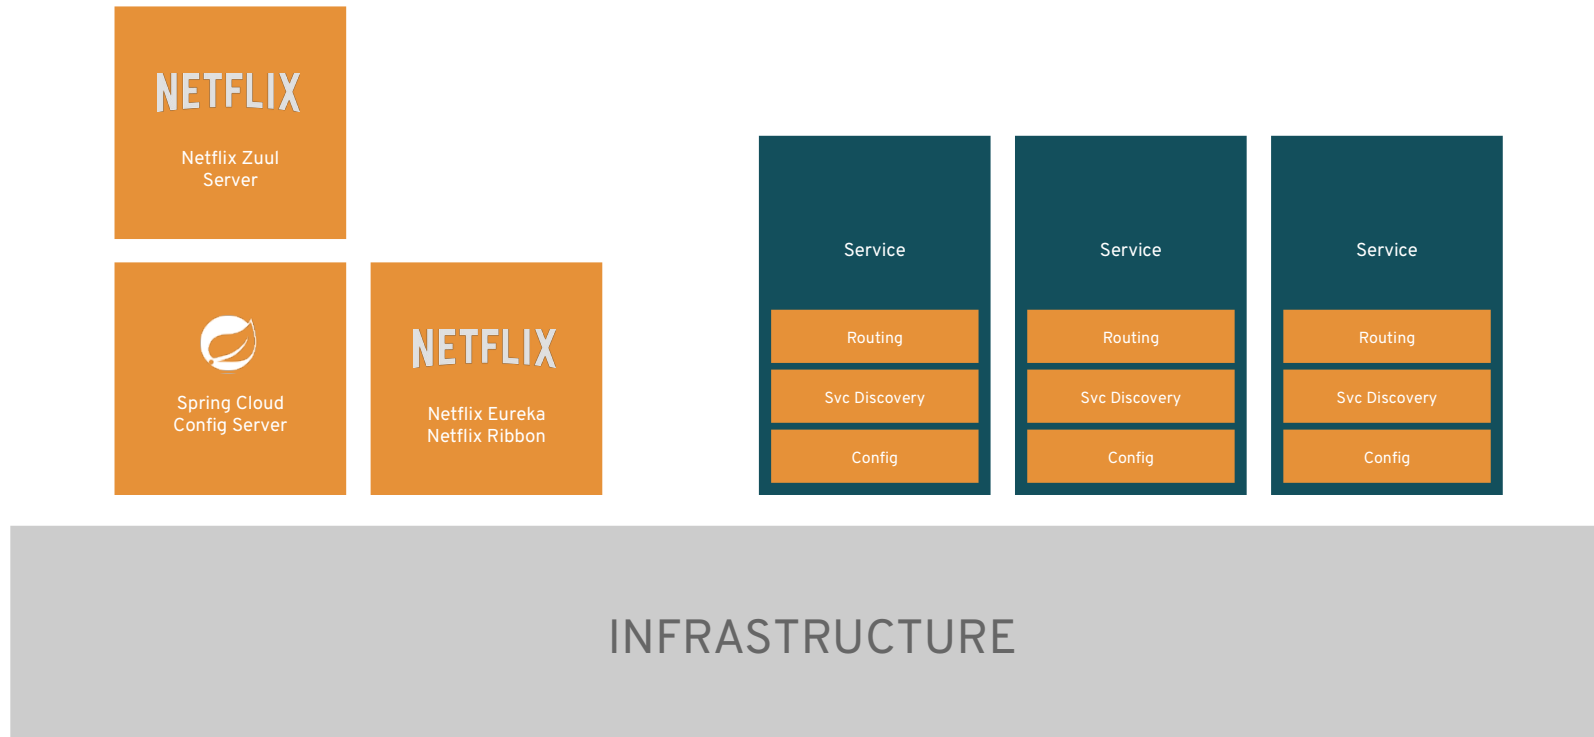
INFRASTRUCTURE

# SERVICE DISCOVERY

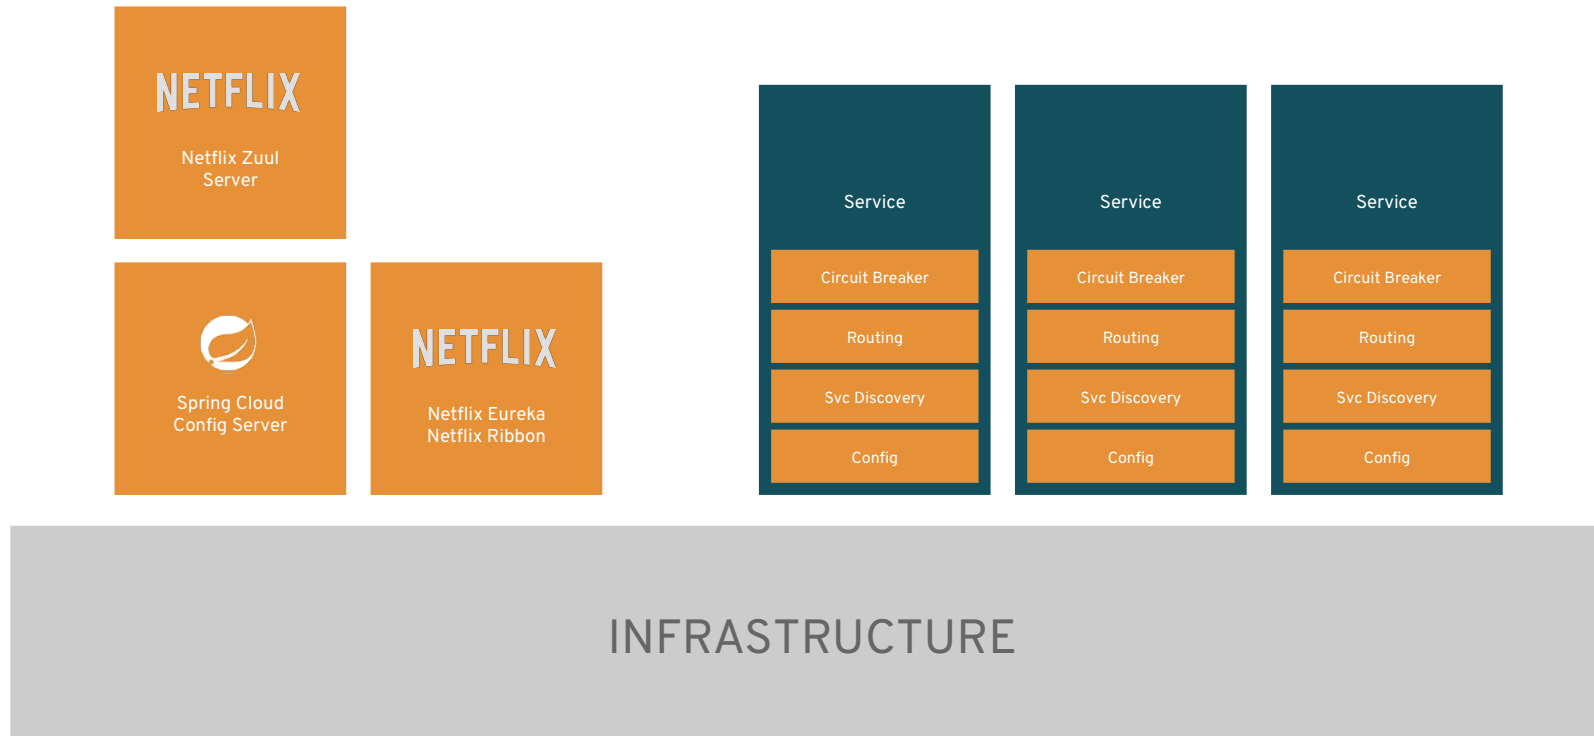


INFRASTRUCTURE

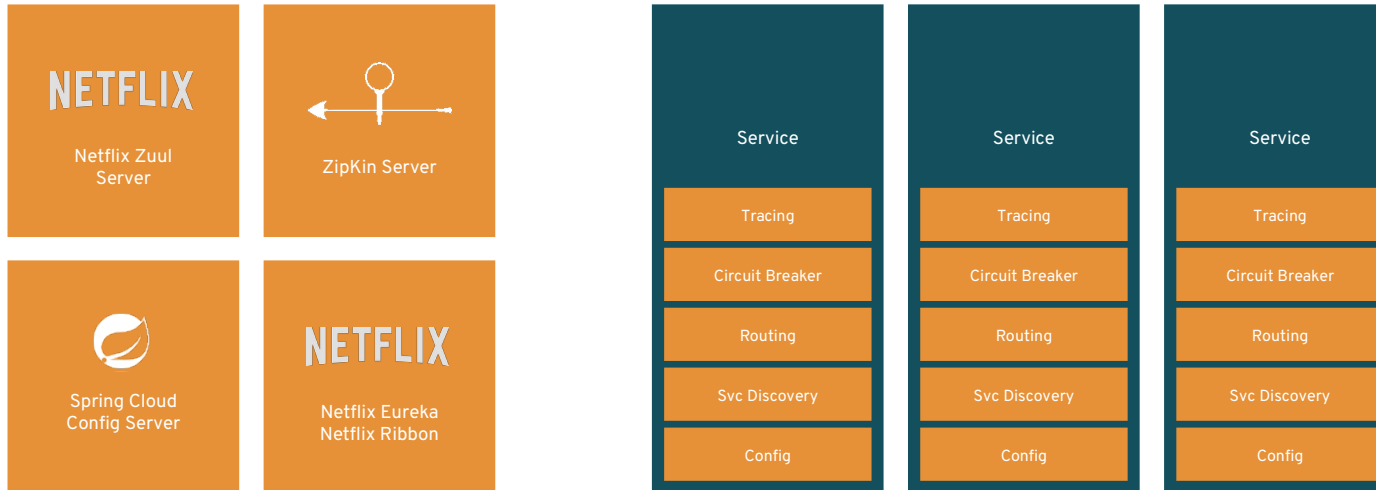
# DYNAMIC ROUTING



# FAULT TOLERANCE



# TRACING AND VISIBILITY



INFRASTRUCTURE

# WHAT ABOUT...?

POLYGLOT  
APPS



EXISTING  
APPS

**THERE SHOULD BE A  
BETTER WAY**



# ADDRESS THE COMPLEXITY IN THE INFRASTRUCTURE

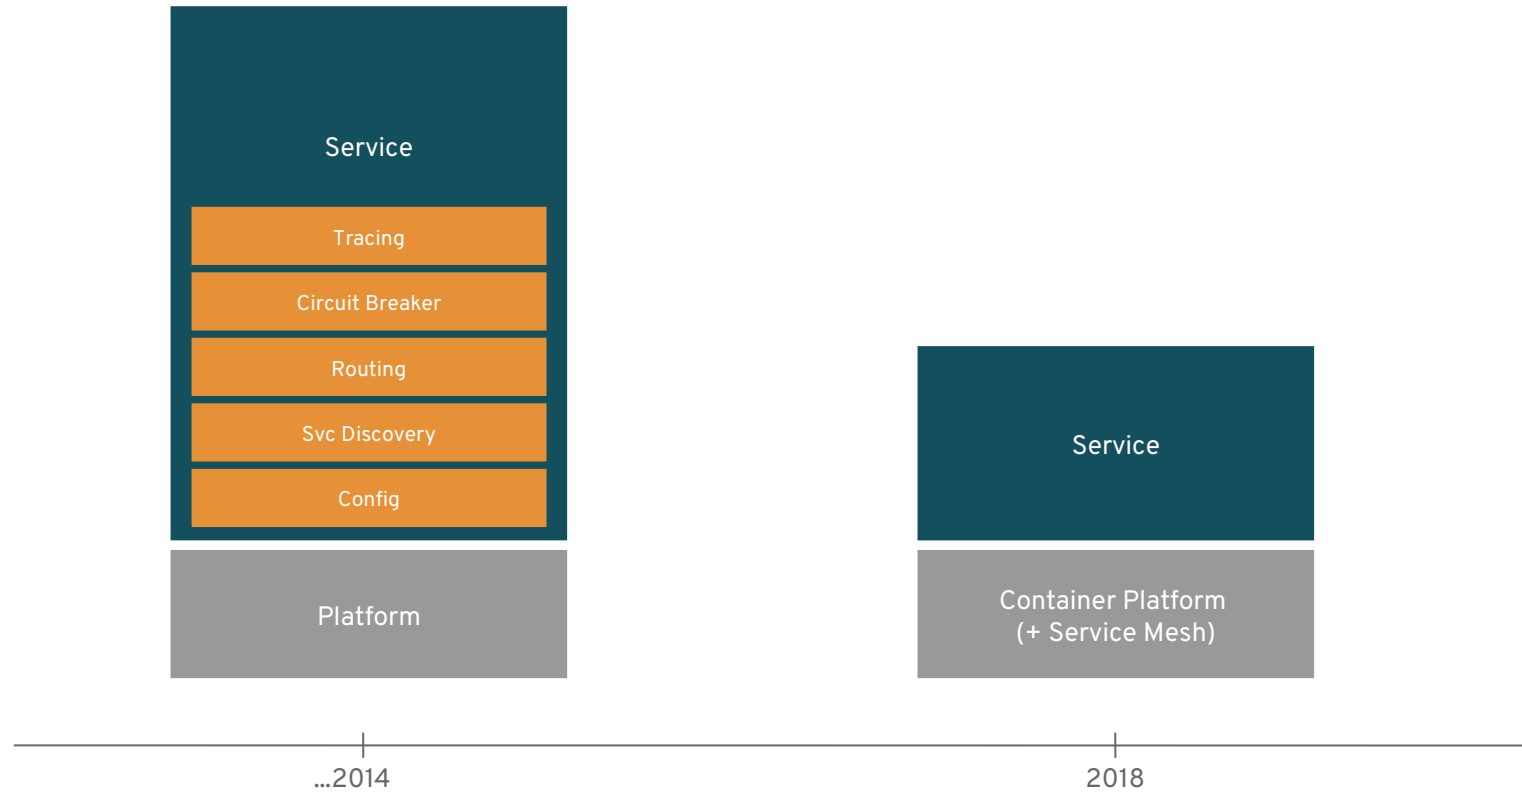


# SERVICE MESH

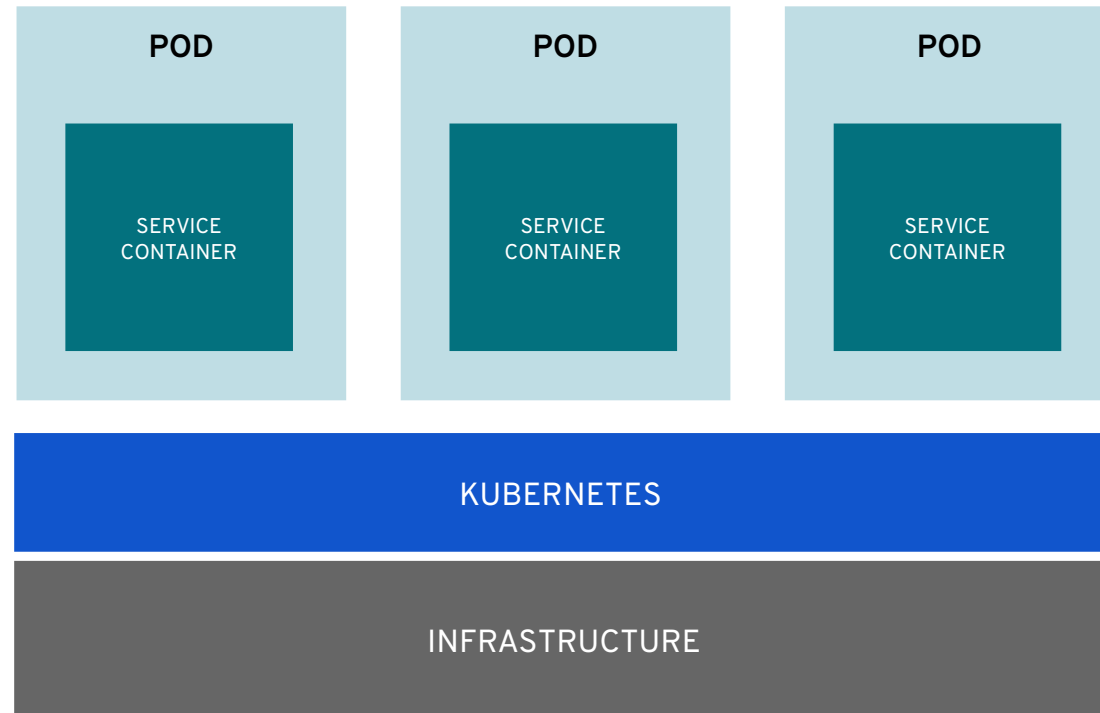
A dedicated infrastructure layer for service-to-service communications



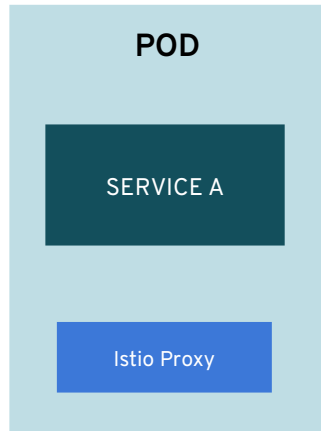
# MICROSERVICES EVOLUTION



# AUTOMATING CONTAINER DEPLOYMENT



# SIDECARS

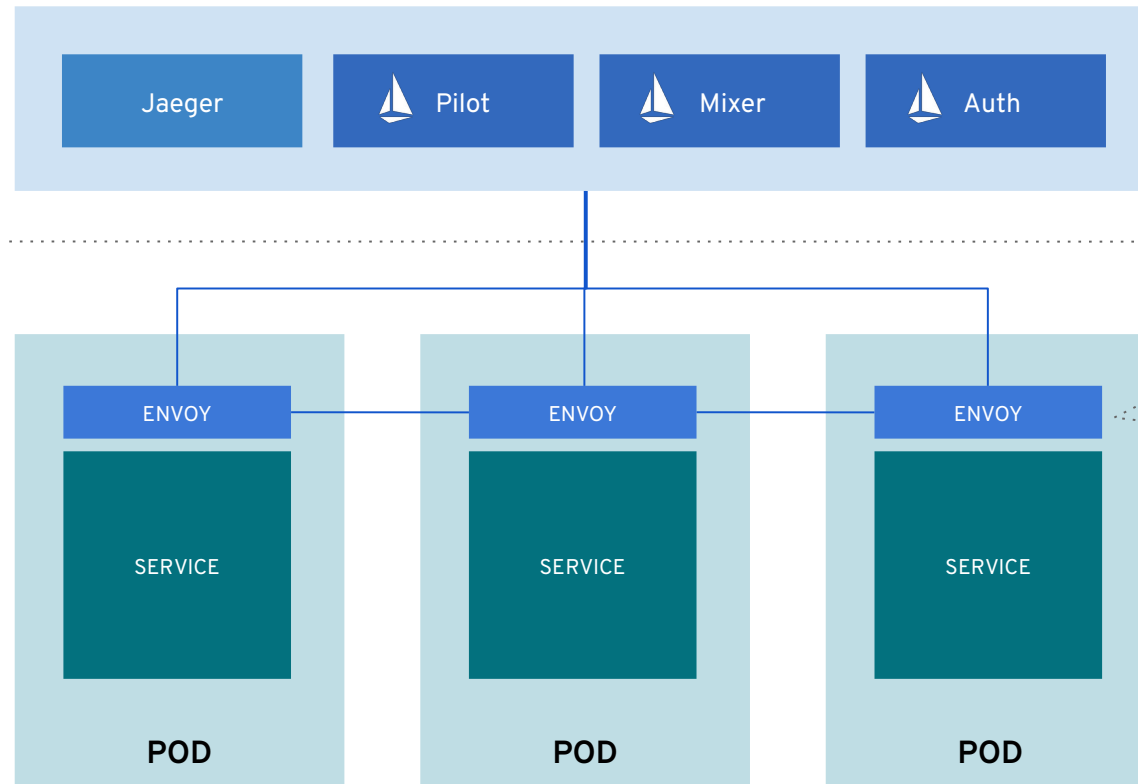


- Two or more containers deployed to same pod
- Share
  - Same
    - Namespace
    - Pod IP
  - Shared lifecycle
- Used to enhance the co-located containers
- Istio Proxy (L7 Proxy)
  - Proxy all network traffic in and out of the app container

Source: <http://blog.kubernetes.io/2015/06/the-distributed-system-toolkit-patterns.html>

# SERVICE MESH ARCHITECTURE

Control Plane



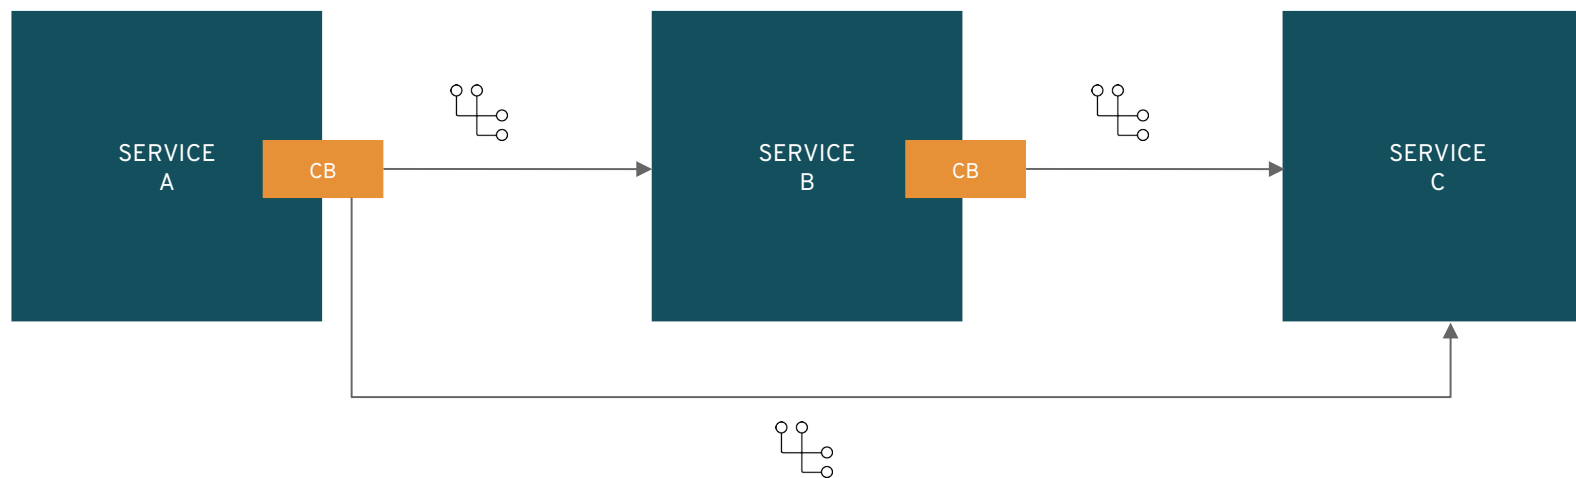
Data Plane

# MAJOR FUNCTIONALITY

# FAULT TOLERANCE

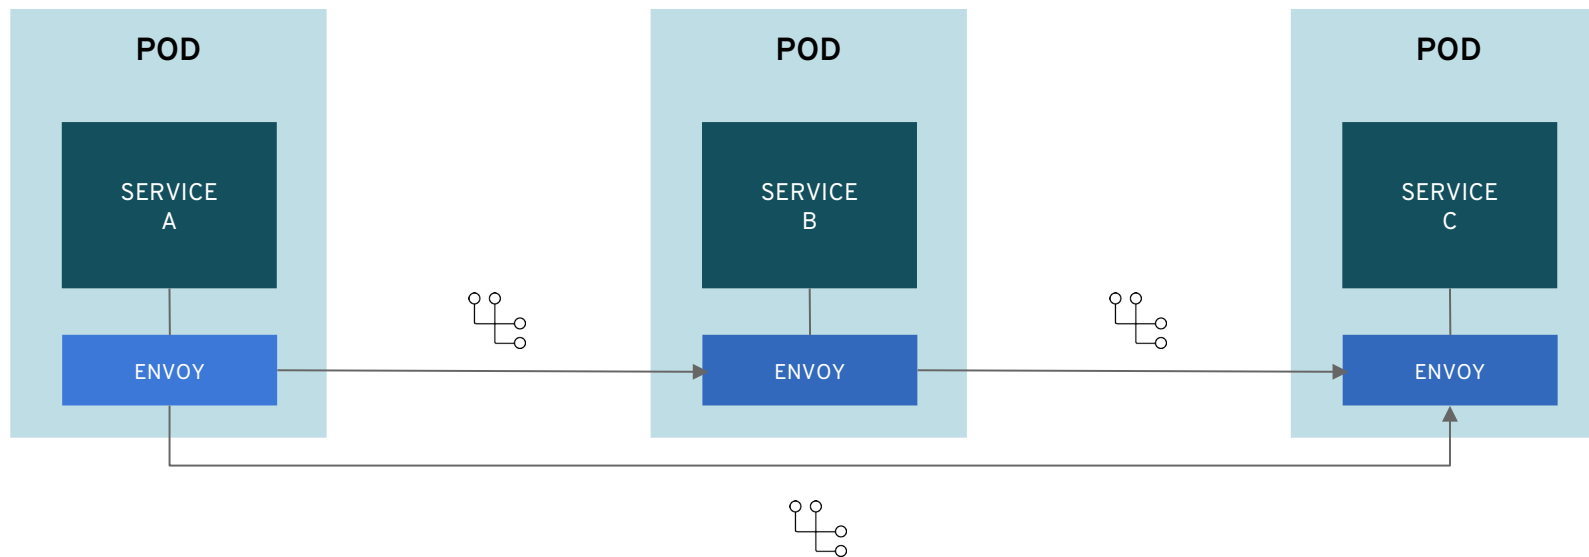


# CIRCUIT BREAKERS **WITHOUT** ISTIO



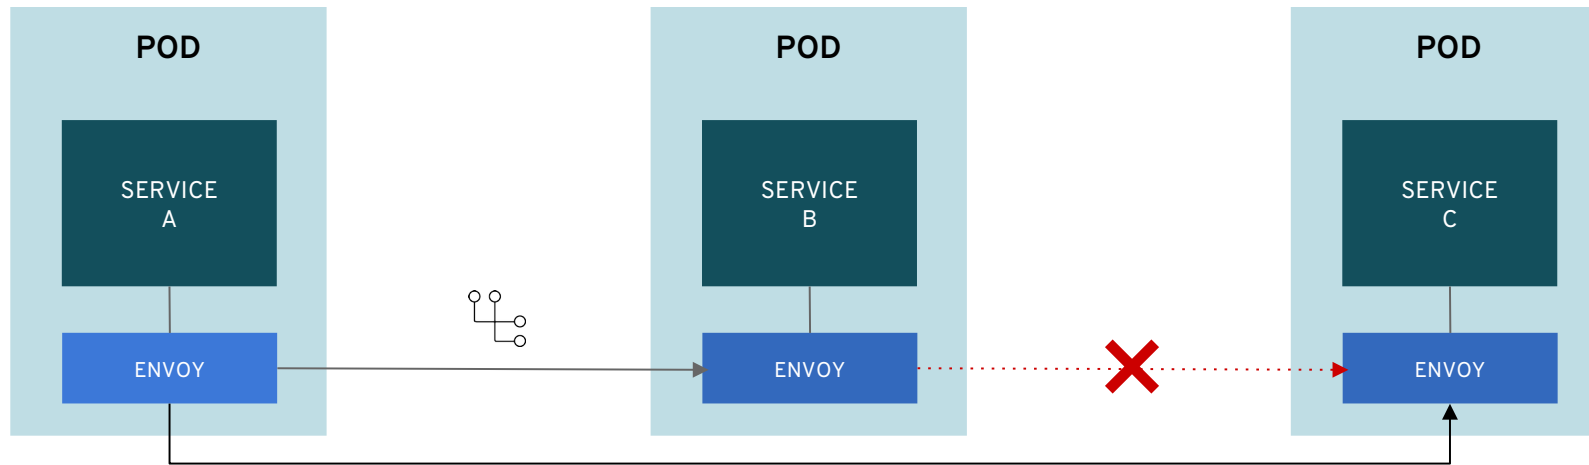
coupled to the service code

# CIRCUIT BREAKERS WITH ISTIO



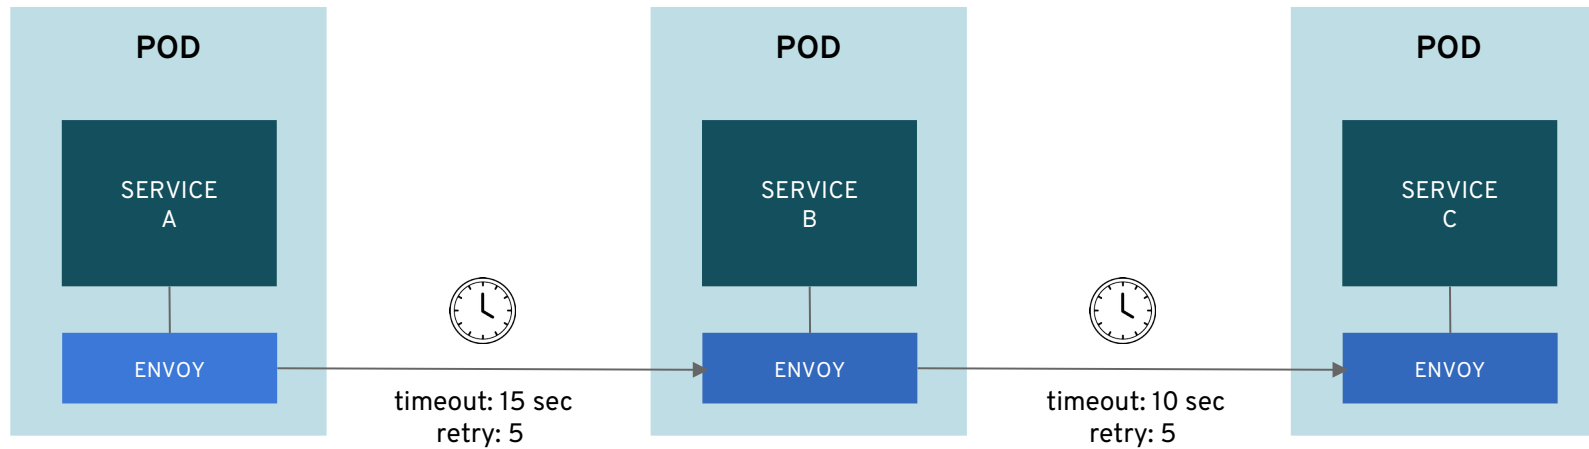
transparent to the services

# CIRCUIT BREAKERS WITH ISTIO



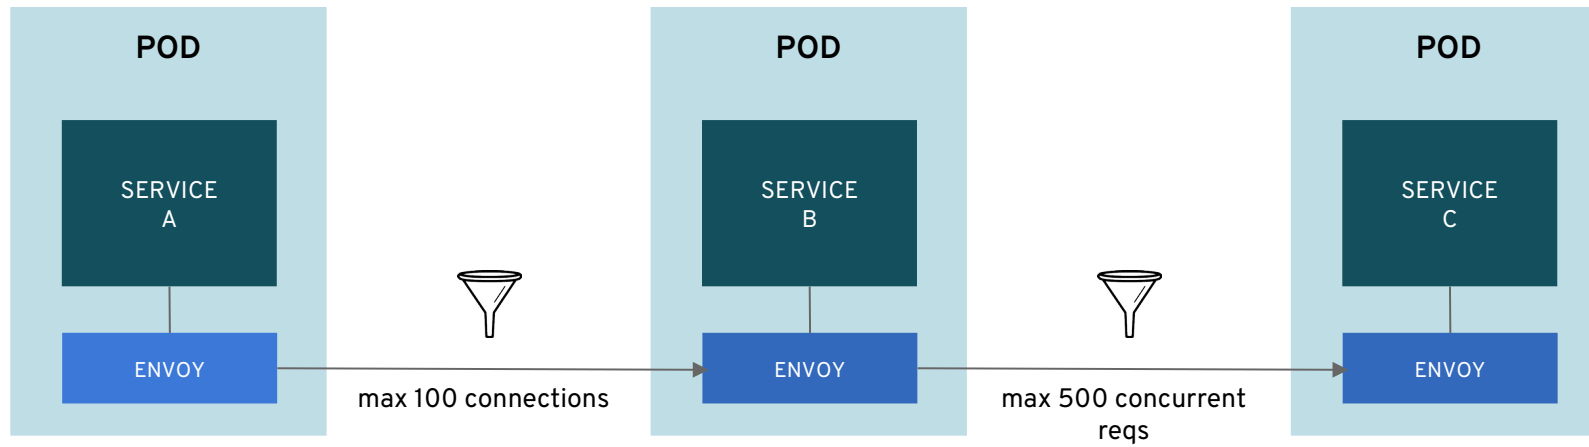
improved response time with global circuit status

# TIMEOUTS AND RETRIES WITH ISTIO



configure timeouts and retries, transparent to the services


# RATE LIMITING WITH ISTIO



limit invocation rates, transparent to the services


# SERVICE SECURITY

# SECURE COMMUNICATION **WITHOUT** ISTIO



coupled to the service code

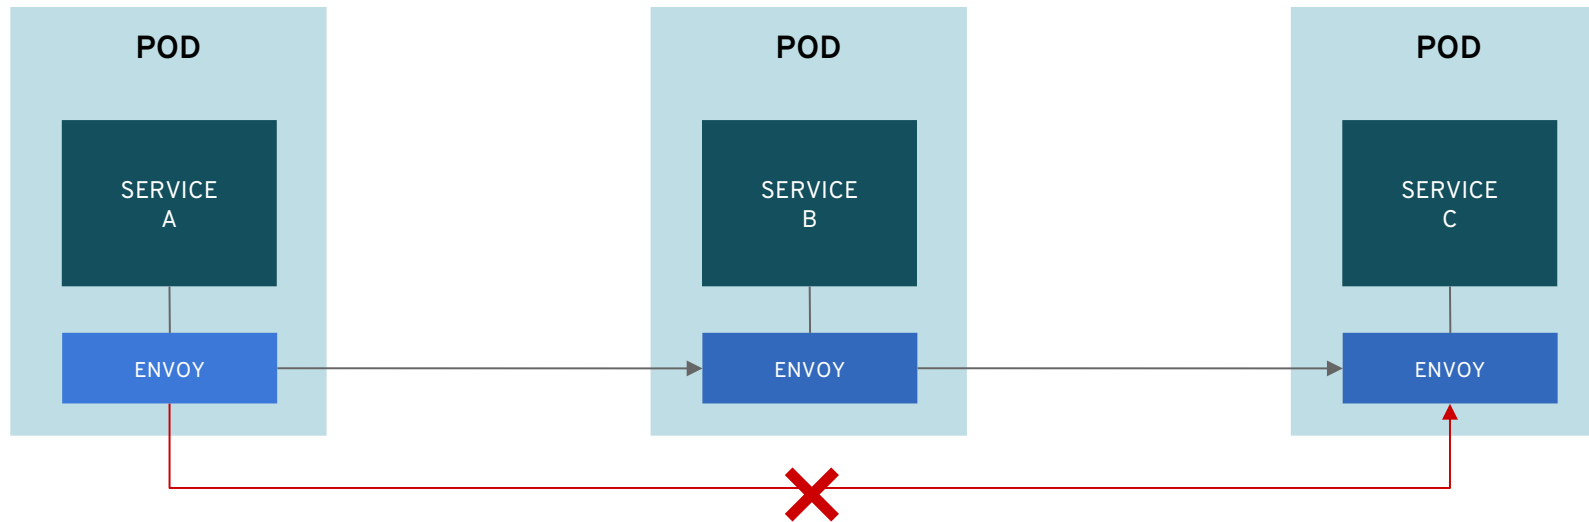
# SECURE COMMUNICATION WITH ISTIO



mutual TLS authentication, transparent to the services



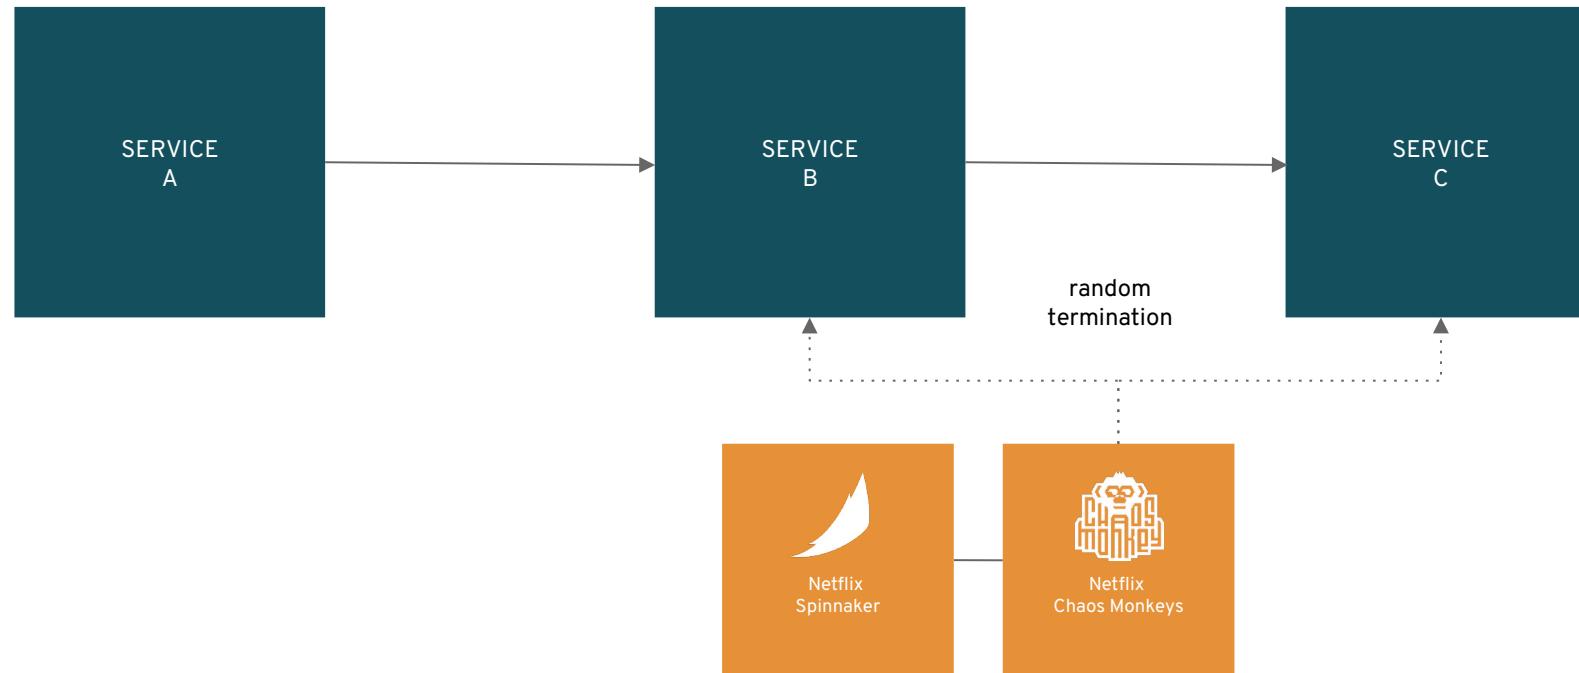
# CONTROL SERVICE ACCESS WITH ISTIO



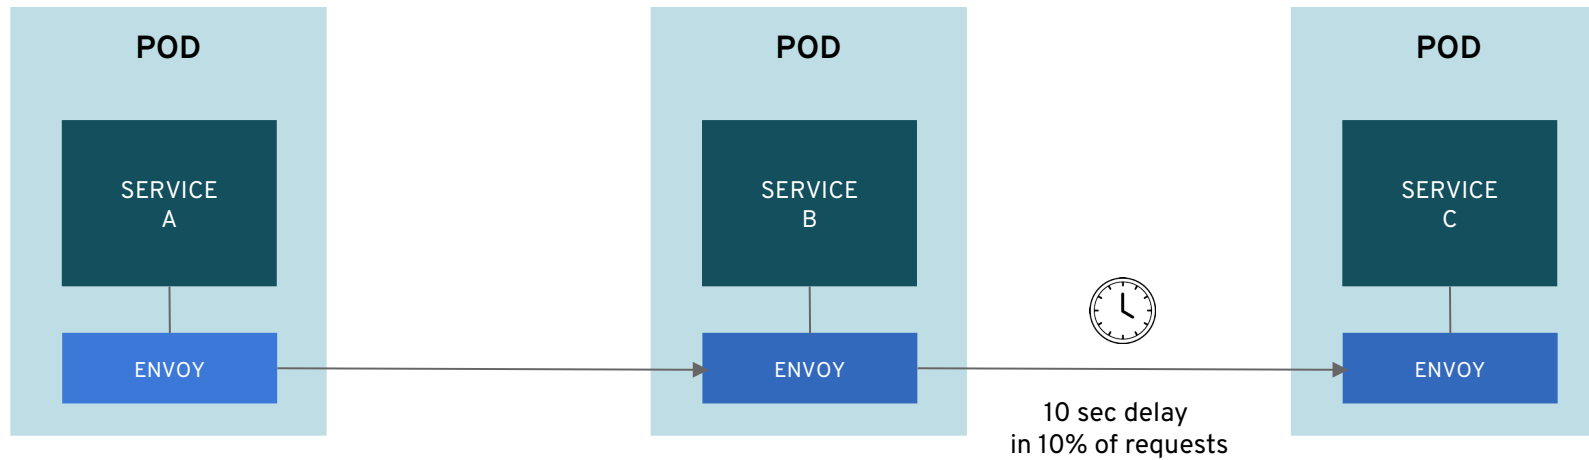
control the service access flow, transparent to the services

# CHAOS ENGINEERING

# CHAOS ENGINEERING WITHOUT ISTIO

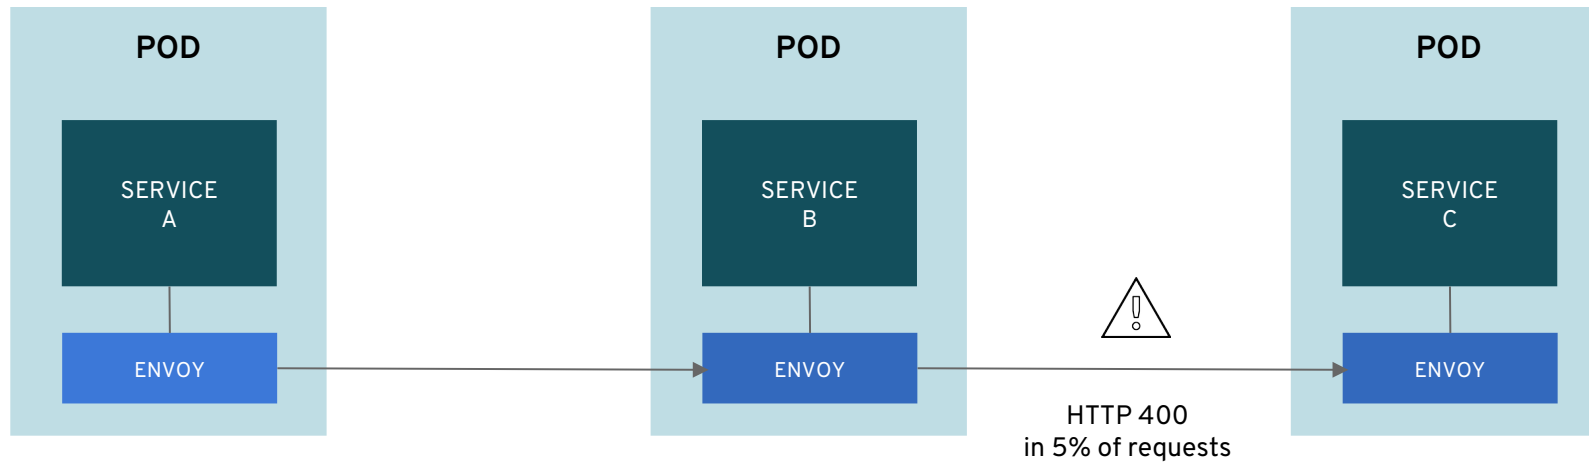


# CHAOS ENGINEERING WITH ISTIO



inject delays, transparent to the services

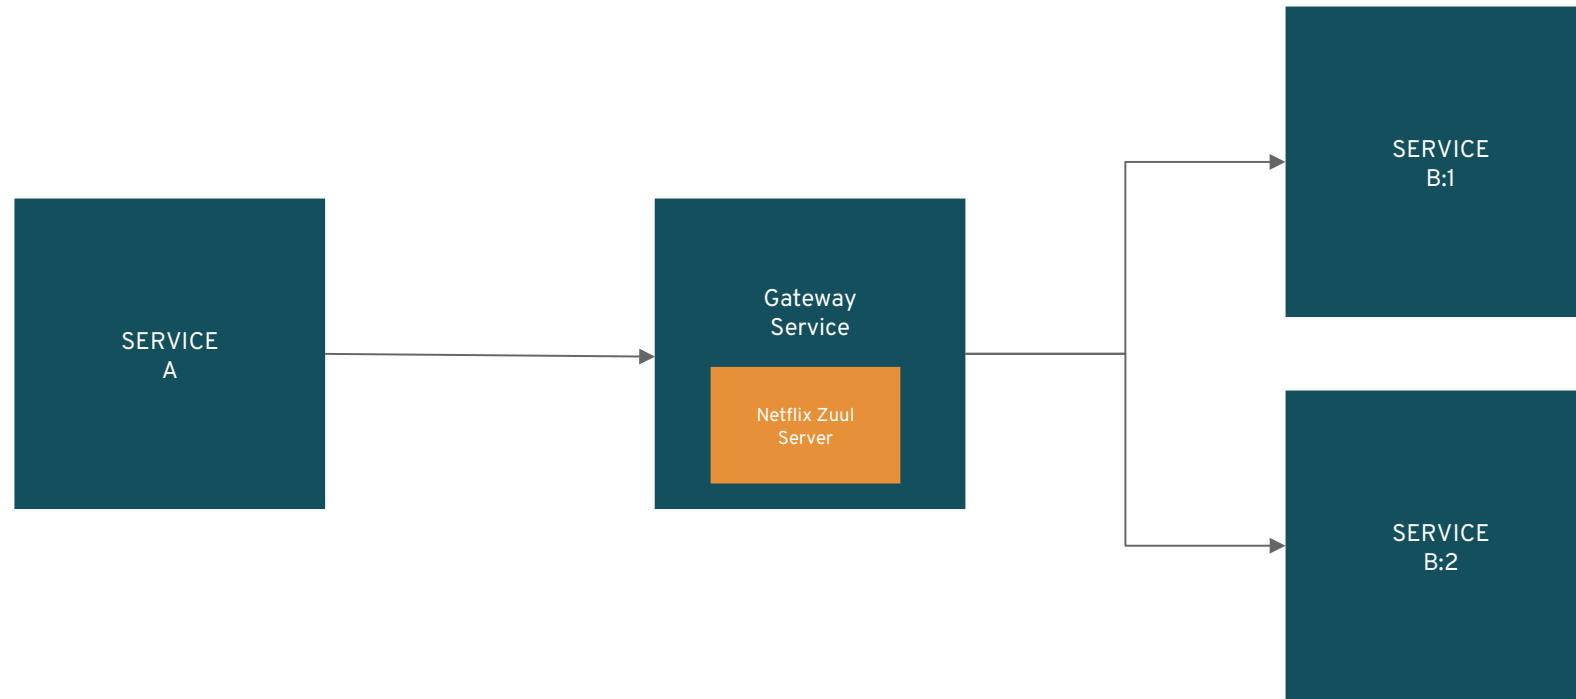
# CHAOS ENGINEERING WITH ISTIO



inject protocol-specific errors, transparent to the services

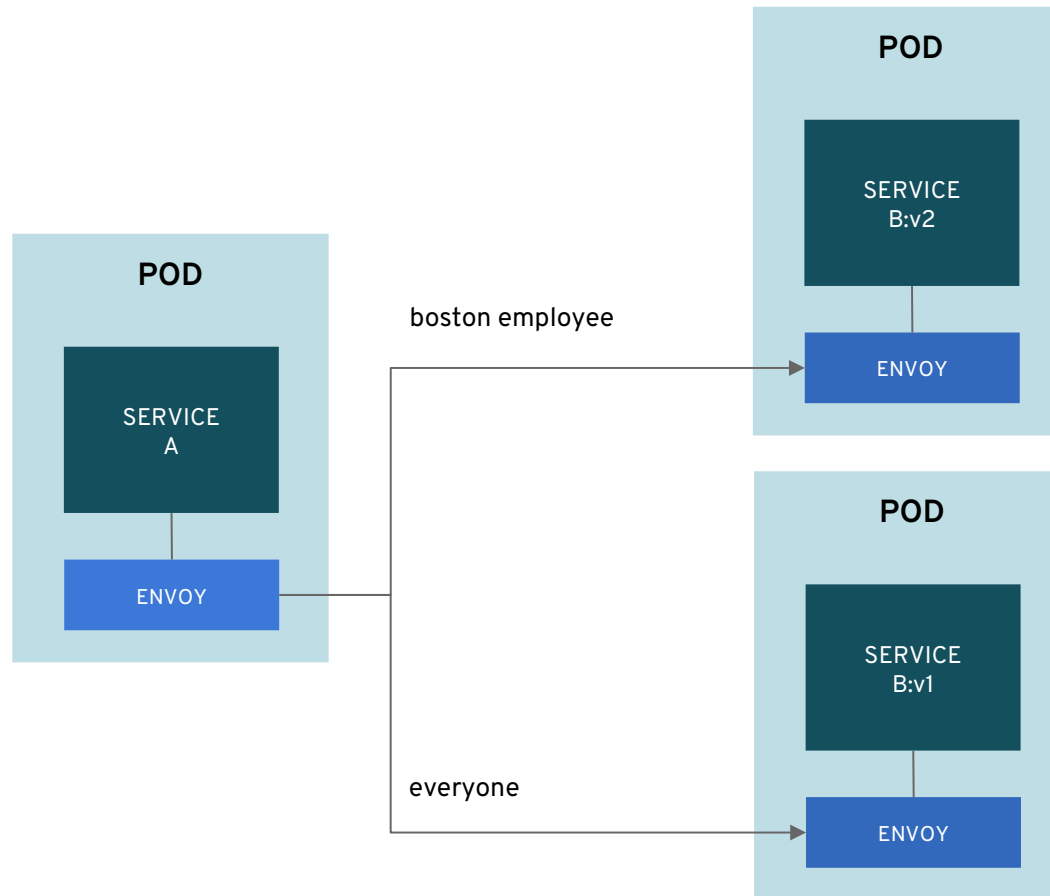
# DYNAMIC ROUTING

# DYNAMIC ROUTING **WITHOUT** ISTIO



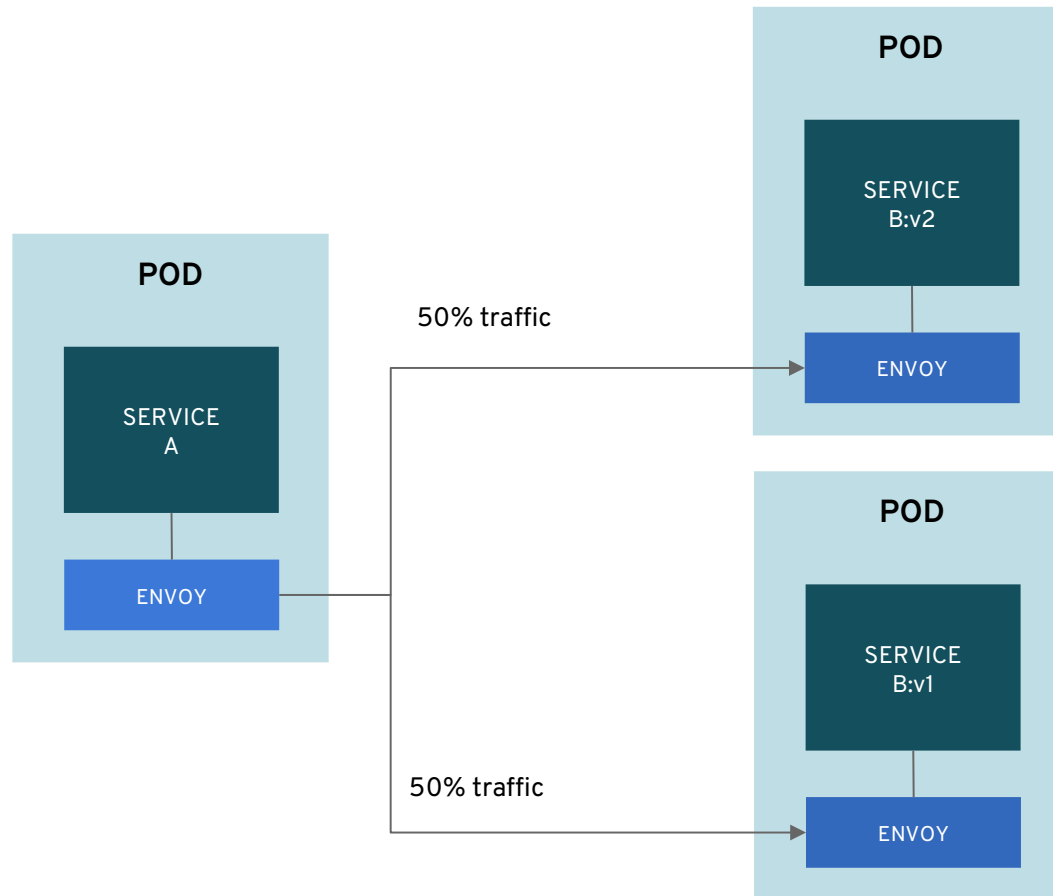
custom code to enable dynamic routing

# CANARY DEPLOYMENT WITH ISTIO

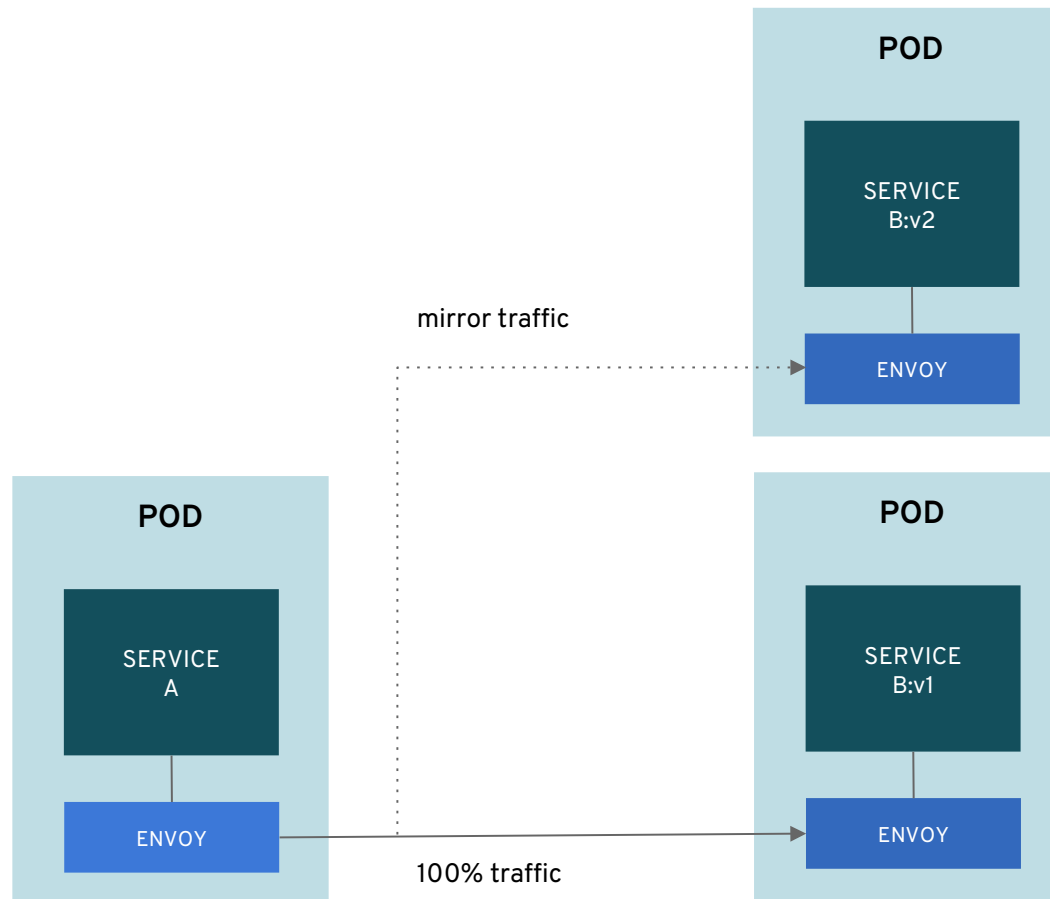




# A/B DEPLOYMENT WITH ISTIO

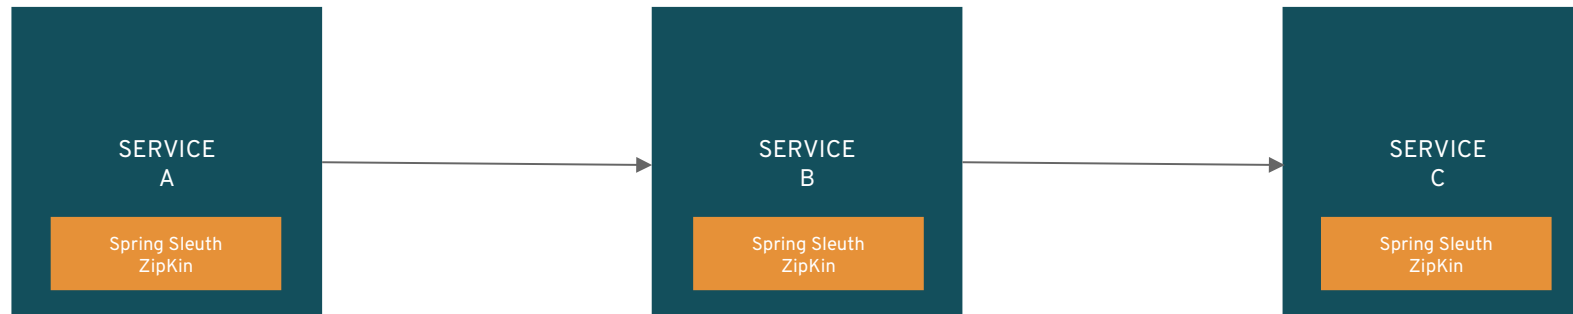


# DARK LAUNCHES WITH ISTIO




# DISTRIBUTED TRACING (JAEGER)

# DISTRIBUTED TRACING **WITHOUT** ISTIO




code to enable dynamic tracing

# DISTRIBUTED TRACING WITH ISTIO & JAEGER



discovers service relationships and process times, transparent to the services



# SERVICE MESH OBSERVABILITY (KIALI)

kiali

Overview

Graph

Applications

Workloads

Services

Istio Config

Distributed Trac...

Namespace: bookinfo

Graph

DisplayEdge LabelsGraph TypeVersioned appFind...xHide...x?

Feb 18, 16:07:37 ... Feb 18, 16:08:37

FetchingLast minEvery 15 sec

```
graph LR; ingress[istio-ingressgateway 1.1.0] --> productpage; subgraph productpage; productpage_v1[v1]; end; productpage_v1 --> details; subgraph reviews; reviews_v1[v1]; reviews_v2[v2]; reviews_v3[v3]; end; productpage_v1 --> reviews_v1; reviews_v1 --> details; reviews_v1 --> reviews_v2; reviews_v1 --> reviews_v3; reviews_v2 --> ratings; reviews_v3 --> ratings; subgraph ratings; ratings_v1[v1]; ratings_v2[v2]; end; reviews_v2 --> ratings_v1; reviews_v3 --> ratings_v2; subgraph details; details_v1[v1]; end; details_v1 --> productpage_v1; subgraph mongodb; mongodb_v1[v1]; end; ratings_v1 --> mongodb_v1; ratings_v2 --> mongodb_v1;
```

Namespace: bookinfo  
applications, services, workloads

Current Graph:  
9 apps  
5 services  
14 edges

HTTP Traffic (requests per second):

Total	%Success	%Error
3.68	100.00	0.00

0255075100%

OK3xx4xx5xx

HTTP - Total Request Traffic min / max:  
RPS: 3.60 / 3.60, %Error 0.00 / 0.00

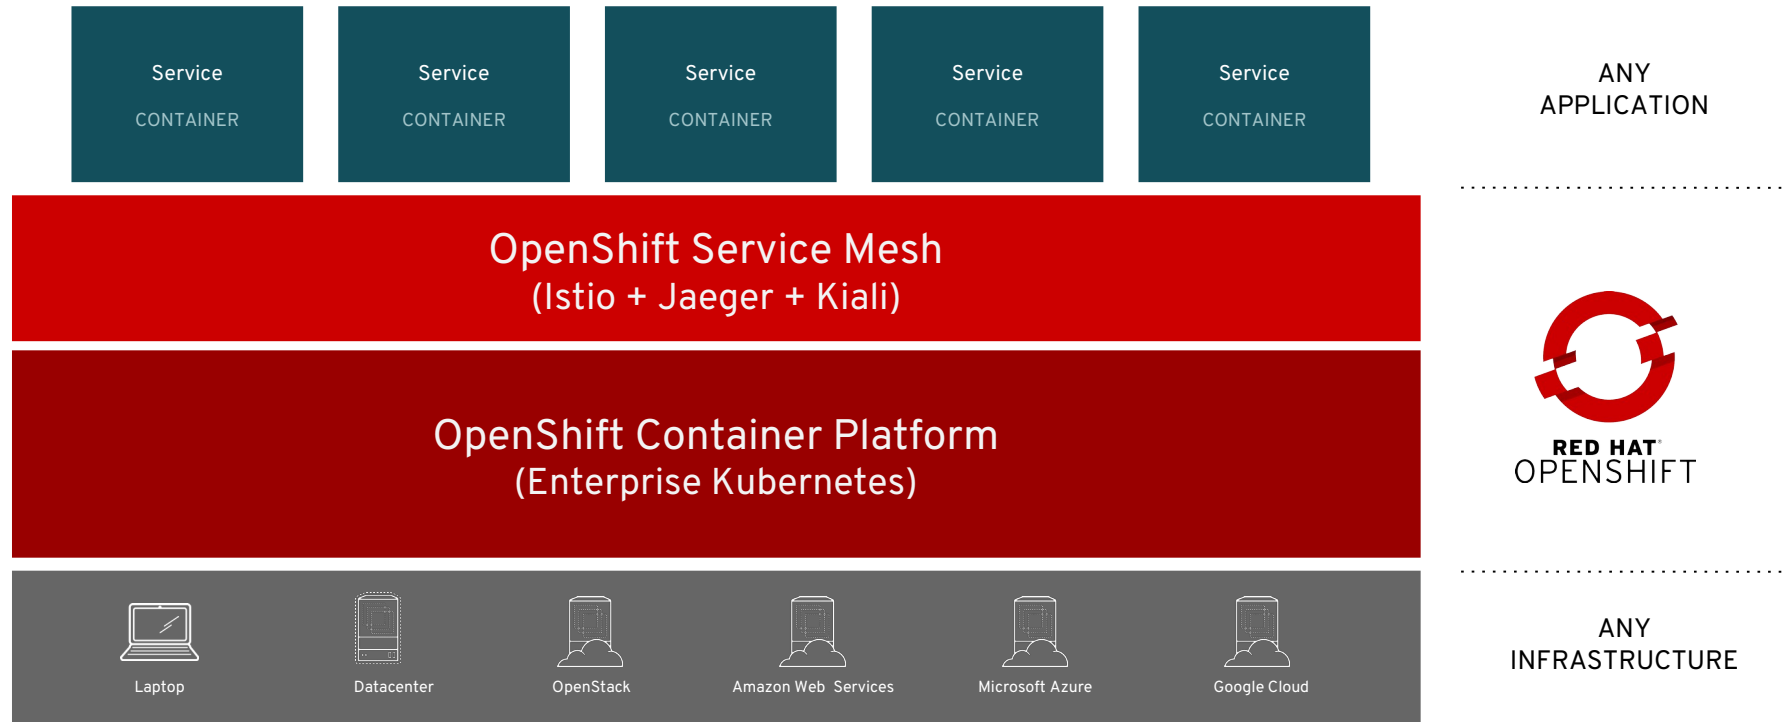
TCP - Total Traffic - min / max:  
Sent: 143.00 / 143.00 B/s  
Received: 115.67 / 115.67 B/s

+ -

12


Legend

# DISTRIBUTED SERVICES PLATFORM





# References



## *How to explain service mesh in plain English*

<https://enterpriseproject.com/article/2019/6/service-mesh-plain-english>

## *OpenShift Commons 2019*

<https://blog.openshift.com/wp-content/uploads/State-of-the-Platform-Services-Integrated-1.pdf>

## *Service Mesh Documentation (OCP 4.2)*

Architecture:

[https://docs.openshift.com/container-platform/4.2/service\\_mesh/service\\_mesh\\_arch/understanding-ssm.html](https://docs.openshift.com/container-platform/4.2/service_mesh/service_mesh_arch/understanding-ssm.html)

Install:

[https://docs.openshift.com/container-platform/4.2/service\\_mesh/service\\_mesh\\_install/preparing-ssm-installation.html](https://docs.openshift.com/container-platform/4.2/service_mesh/service_mesh_install/preparing-ssm-installation.html)

Day 2:

[https://docs.openshift.com/container-platform/4.2/service\\_mesh/service\\_mesh\\_day\\_two/prepare-to-deploy-applications-ssm.html](https://docs.openshift.com/container-platform/4.2/service_mesh/service_mesh_day_two/prepare-to-deploy-applications-ssm.html)

# Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.



[linkedin.com/company/red-hat](https://linkedin.com/company/red-hat)



[youtube.com/user/RedHatVideos](https://youtube.com/user/RedHatVideos)



[facebook.com/redhatinc](https://facebook.com/redhatinc)



[twitter.com/RedHat](https://twitter.com/RedHat)